NASA-TM-109210

FI Lib

51-60

17663

1 11

N93-72606

Comparisons Between Intel 386 And i486 Microprocessors

Yuan-Kwei I

(NASA-TM-109210) [PROCESSOR TECHNOLOGY STUDIES] (NASA) 47 p

N93-72605
--THRU--
N93-72608
Unclas

Z9/60 0176580

PRELIMINARY REPORT

September 1989

NASA

Ames Research Center

Moffett Field, California 94035

COMPARISONS BETWEEN INTEL 386 AND i486 MICROPROCESSORS

1. THE i486 KEY ELEMENTS

(1) Full 386 compatibility: The i486 includes a highly-optimized integer unit, which is fully compatible with the 386.

(2) On-chip floating point unit: The i486 contains a FPU, which is compatible with the 387 numeric coprocessor. The 386 uses a separate 387 coprocessor.

(3) On-chip 8-Kbyte cache: The i486 contains 8-Kbyte cache, which can handle both instruction code and data.

(4) Built-in multiprocessor support: The i486 supports multiprocessor instructions, cache consistency protocols, second level cache, and other multiprocessor support hooks.

(5) Faster CPU: Intel claims that the initial 25-MHz version of the i486 will run between 2 and 4 times faster than a 25-MHz 386. The 33-MHz i486 will be announced later. This seems to be supported by the benchmark results reported by October, 1989 MIPS magazine (See Appendix 4, Table 2, Page 43).

(6) Faster bus: The i486 bus is significantly faster than the 386 bus. Differences include a 1X clock for three heavily used instructions (discussed in Performance Section), parity support, burst cycles, and 8-bit bus support.

2. COMPARISON OF INSTRCTION SET ARCHITECTURE

The i486 implements essentially the same instruction set, programming model, memory organization, process organization, process management facilities, and so forth as the original 386, though the i486 significantly outperforms the 386. Intel made no significant changes to the 386 ISA except for the addition of three new instructions required for cache support (INVLPG, INVD, and WBINVD) and three for multiprocessing functions (CMPXCHG, XADD, and BSWAP).

3. THE i486 ON-CHIP CACHE CHARACTERISTICS

The on-chip cache is fundamental to achieving the high performance level of the i486. Though the 8-Kbyte cache is considerably smaller than the 32-Kbyte external caches that are built in many 386-based PCs, it is considerably more sophisticated. Wide buses from the cache let the instruction prefetcher grab blocks of code 16 bytes at a time, instead of four. In contrast to most processors with on-chip caches, the i486 cache is unified to hold both code and data. This generally provides more efficient cache utilization than would, for example, separate 4-Kbyte code and data caches. It also solves, quite cleanly, problems that might otherwise arise in executing application programs that include self-modifying code.

The most often cited disadvantage of a unified cache design is that code and data fetches can collide, stalling instruction execution. The i486 prefetch system avoids this hazard in two ways. Prefetching instructions 16 bytes at a time reduces the frequency of collisions, and when collisions do occur, the data request is serviced first. The execution unit can generally keep busy finishing instructions already in the prefetch queue.

4. THE i486 MULTIPROCESSOR SUPPORT

The 486's on-chip cache uses a write-through memory update policy: all writes are passed through to main memory. As a result, the 486 produces considerable bus activity, even though most read cycles are satisfied by the cache. The on-chip snooping functions will provide the needed cache invalidations when DMA writes to cached memory locations. Because of the small cache size and write-through memory policy, shared-memory multiprocessor systems will require second-level caches. Implementing such caches with write-back support is a complex task, but will be greatly simplified when Intel's cache controller becomes available next year. Second-level caches will also be desirable on the 33-MHz i486 uniprocessor systems.

5. COMPARISON OF PERFORMANCE

Intel claims that performance of the i486 at 25 MHz will be about 37,000 Dhrystones, 6.1 M Whetstones (double precision), or about 15-20 VAX MIPS. Intel also claims the performance enhancement of 2 to 4 times over the 386. The fully-redesigned integer unit and on-chip cache reduce the average number of clocks per instruction by a factor of about 2.5. The on-chip FPU executes floating-point instructions about three times faster than the 386/387 combination. And while the initial i486 clock speed is 25-MHz, the part was designed to scale to clock rates higher than those achievable with the 386. The comparisons of instruction execution time between the 386 and i486 are shown in Table 1.

Instruction	Optimum Clock Cycle Count	
	386/387	i486
Register-to-register add	2	1
Memory load (16 bit)	2	1
Memory store (16 bit)	2	1
Memory-to-register add (16 bit)	6	2
Register-to-memory add (16 bit)	7	3
Unconditional jump	8	3
Call	8	3
Return	11	5
Floating-point load (64 bit)	25	3
Floating-point add (80 bit)	23-31	8-20
Floating-point multiply (80 bit)	29-57	16
Floating-point divide (80 bit)	88	73

Table 1. Instruction execution time comparisons

6. COMPARISON OF POWER CONSUMPTION

Microprocessor	Power Dissipation
386	3W*
387	1.5W
i486 (25-MHz)	5W
i486 (33-MHz)	6W

Table 2. Power consumption comparisons

* Estimated from power supply current for the 25-MHz 386

As mentioned before, the i486 contains both integer unit (compatible with the 386) and floating-point unit (compatible with the 387). The 25-MHz i486 power dissipation (5W) is 0.5W higher than the sum of the 386 and 387 (4.5W), but the i486's 5W includes on-chip 8-Kbyte cache. Therefore, the strict power consumption requirement for the Space Station Freedom Data Management System (DMS) may not cause a problem for using the i486, instead of the 386.

7. COMPARISON OF RADIATION HARDENING POTENTIAL

The 386 uses CHMOS III and CHMOS IV process technologies. The 387 uses CHMOS III technology. The i486 uses CHMOS IV technology. The "386 Radiation Analysis" based on CHMOS IV was studied and presented by J.W. Devlin on 2/2/89 in the DMS Working Group meeting. Since the i486 uses the same kind of process technology, the strict DMS radiation hardened requirement may not cause a problem for using the i486, instead of the 386. In the meeting that we had with Intel microprocessor development managers, Intel expressed their interest in qualifying the i486 for the military and space applications.

8. RECOMMENDATION FOR DMS

The i486 demonstrates the following clear advantages over the 386:

(1) PERFORMANCE: Execution speed is important for the time-critical, real-time processing environment, especially for DMS operating system/Ada run time environment (OS/Ada RTE), network operating system (NOS), and standard services.

(2) LONGER LIFE CYCLE: The i486 can take advantage of the vast software that have been written for the 386 because it is fully upward compatible. But a new software written for the i486 may not be able to run on the 386 because it may not be downward compatible.

(3) MULTIPROCESSOR SUPPORT: Multiprocessor systems will be supported by the new Mach-influenced version of Unix System V, Release 4, which Intel, AT&T, and others are working on. This upcoming version of Unix should support mixed systems with multiple i486 and 860/960 processors. IBM AIX operating system will be influenced by this industrial trend. In the meeting that we had with Intel microprocessor development managers, we were informed that the next generation Intel microprocessor is projected to have a significantly larger cache and the write-through and write-back protocols to support a 4 CPUs multiprocessor system. Ada tasking can really benefit from this multiprocessor feature when it is available.

Work Package 2 (WP-2) proposed buying IBM PS/2 Model 80 for simulating DMS Standard Data Processor (SDP). The Model 80 can not be upgraded to i486. IBM PS/2 Model 70-A21 is recommended because it is the only PS/2 model that can be upgraded to i486 as of September 1989.

These comparisons between 386 and i486 were made on a quick and preliminary analysis. An in-depth study to evaluate the risk and advantage is recommended, especially on the power consumption, radiation hardened issues, and the quality and reliability concerns for a new product.

9. APPENDIX

(1) A page on 386 features and architecture from Intel "Microprocessor and Peripheral Handbook".

(2) A page on i486 features and architecture from Intel "i486 Microprocessor" book.

(3) Article on "The Intel 80486" in the June 1989 issue of the "MIPS" magazine.

(4) Article on "486 Times Two" in the October 1989 issue of the "MIPS" magazine.

Note: The Item 1 and 2 books are free and can be ordered by calling Ms. Ann Schulz, Intel Field Sales Engineer, at (408) 970-1649.

N93-72407

Current Processor Technology Survey

Yuan-Kwei Liu

June 1990

National Aeronautics and
Space Administration

TABLE OF CONTENTS

1. INTRODUCTION

This report provides a technology survey on current commercially available processors. This report is part of an effort funded by OAET and the Space Station Freedom (SSF) Advanced Development program to determine computational needs and capability for advanced automation and robotics. It will be used in conjunction with a user requirements survey to determine a plan to meet computational requirements for future NASA missions.

Section 2 is a glossary for this report for easy reference to commonly used technical terms. The section consists of five categories: system design, cache, general architecture, performance, and military standards.

The domains of the microprocessors or architectures discussed in this report include: (1) those with significant space or military applications; (2) those have been used widely in the commercial market; (3) new processors or architectures.

According to these domains, Section 3 covers Intel 386DX, 386SX, i486, i860, and 80960 processors. Section 4 discusses MIPS R2000, R3000, and R6000 processors. Section 5 discusses the Sun SPARC architecture. Section 6 covers the Motorola MC68030, MC68040 and MC88000 family of processors. Section 7 discusses the IBM RS/6000 processor. Please notice that the RS/6000 is a proprietary design of IBM and is not related to the MIPS R6000.

Section 8 discusses Japan's TRON project and microprocessor, which are being proposed by NASDA (National Space Development Agency of Japan) for use in the SSF Japanese Experiment Module (JEM).

Section 9 covers the 16-bit MIL-STD-1750A instruction set architecture, its limitations, and the Radiation-Hardened 32-Bit Processor (RH-32) program to replace the 1750A standard.

Section 10 compares different processors or architectures which belong to the same category or have similar capability. Intel 386/387DX and i486 are compared due to the possibility of replacing the proposed 386/387DX with the i486 in the SSF. The RISC-based Intel i860, MIPS R3000, Sun SPARC and Motorola MC88000 are also compared.

Section 11 predicts some possible future trends of microprocessors.

2. DEFINITION OF TERMS

2-1 System Design

CISC: A Complex Instruction Set Computer, as compared with the RISC-based processors. Intel 386DX and i486, Motorola MC68030 and MC68040, and Digital Equipment Corporation's VAX architecture are normally classified as CISC-based processors.

RISC: A Reduced Instruction Set Computer. While simpler instruction sets are one aspect of the design approach that has been termed RISC, just counting the number of instructions isn't a useful way to characterize an architecture. The key characteristics of RISC design include: (1) simple fixed-format and fixed-length instructions; (2) simple, consistent instruction encoding; (3) (4) delayed branch instruction; (5) most instructions executed in a single clock cycle; (6) large register set (at least 32 registers) or register windows; (7) load/store architecture, where only load and store instructions access memory, and arithmetic and logic instructions operate only on registers. According to the above characteristics, the RISC-based chips include Intel i860, Motorola 88000, MIPS processors, Sun SPARC, IBM RS/6000, etc.

2-2 Cache

Cache: A small, high-speed memory that is inserted between the main memory and the CPU. It holds part of the currently active portion of a program and its data. A cache usually holds between 1 Kbyte to 512 Kbytes of data.

Cache coherency: In multiprocessing systems with each CPU containing a local cache, multiple copies of program variables can exist in the system. Each CPU can be attempting to access and/or modify its copy of the program variables simultaneously. This protocol ensures that only one logically correct value exists for each program variable [1].

Write-back cache (also called copy-back cache): In this cache update policy, only the cache location is updated, the main memory location retains its previous value. Later, when the block containing this updated line is removed from the cache for a new block, the main memory location of the line is updated. The advantage of this policy (compared to write-through policy) is that it potentially generates fewer main memory updates. The disadvantage is that it is likely to cause bursts in memory traffic since a context switch will cause a flood of memory writes [2].

Write-through cache (also called store-through cache): In this cache update policy, the main memory location and the cache location are both updated. The advantages of this policy (comparing to write-back policy) are that main memory is always up to date and the logic needed is simplified. The disadvantage of the policy is increased bus traffic.

2-3 General Architecture

Big endian and little endian: These byte ordering schemes define the order of bytes in multiple-byte data. In the big endian scheme, byte 0 is always the most significant (leftmost) byte. Motorola MC68000, Sun SPARC, and IBM 370 use this scheme. In the little endian scheme, byte 0 is always the least significant (rightmost) byte. Intel 80X86 and DEC VAX use this scheme. Some processors, such as MIPS and Motorola MC88000 support both schemes. Selection of the schemes is important for data compatibility.

Register window: In this register organization scheme, a group of register windows (each with, for example, 32 registers) is arranged as a circular buffer. Each procedure call results in a new window of registers transparently allocated to the new procedure, thereby eliminating the need to save registers on each procedure call. Similarly, as each procedure completes and returns, the system readjusts the current window back to the correct window. Sun SPARC architecture has implemented the register window scheme.

Superscalar technique: Superscalar implementation normally has a Long Instruction Word (LIW) architecture to carry multiple instructions. An idealized superscalar processor contains a pool of execution units, each able to interpret a general-purpose instruction set, and each with access to a common register file, bus interface, and so forth. Such a processor potentially can execute many instructions each clock cycle, and increase the performance [3]. The Intel 80960 family of embedded processors and the IBM RS/6000 have superscalar implementations.

Z-buffer: It is a dynamic, RAM-resident buffer to support graphics processing. It has a one-to-one correspondence with a frame buffer. That is, each pixel in the frame buffer has a corresponding location in the z-buffer. Each z-buffer location contains the depth (z) value (the pixel's distance from the viewer) of the object being displayed at the corresponding pixel in the frame buffer. Before drawing a pixel, the graphics unit compares the z value of the object, at that pixel, with the value in the z-buffer. The unit updates the pixel only if the object is closer to the viewer than that currently in the frame buffer [1]. Intel i860 has implemented the z-buffer.

2-4 Performance

MIPS: This term has three meanings in this report. The first one is an acronym name for "million instructions per second". This term is commonly used to measure a processor or a computer's performance. This MIPS is the clock frequency (in MHz) divided by the number of cycles per instruction. In practice, this term is used to compare performance with a VAX 11/780's performance, which is normally referred as 1 MIPS. The second term stands for MIPS Computer Systems, Inc. This report uses MIPS Computer when referring to the company. The third stands for "Microprocessor without Interlocked Pipeline Stages", which is mentioned in the MIPS R2000/R3000 section.

2-5 Military Standards

MIL-STD-883C: This standard specifies the test methods and procedures for microelectronics in the military application [4]. There are two classes defined in this standard: Class B (for military classification) and Class S (for space classification). Generally speaking, a microprocessor vendor introduces a processor for the commercial market first because of the market size. After a period of time, the vendor may qualify the processor for Class B application, which has less market potential than the commercial market. Finally, the chip may be qualified for Class S application, whose market size is the smallest. The time needed from commercial to Class B normally is longer than from Class B to Class S. The "C" in the 883C is the revision number.

MIL-STD-1750A: This military standard defines a 16-bit computer instruction set architecture (ISA) based on the DEC PDP-11 ISA. The standard may be replaced by the Radiation-Hardened 32-Bit Processor (RH-32) program. Please refer to the section on MIL-STD-1750A Architecture and RH-32 Program for details.

3. INTEL PROCESSORS

3-1. Intel 386DX and 386SX Microprocessors

The 80386 processor is called 386DX now by Intel to distinguish it from the 386SX. The 386DX is the baseline processor for the Space Station Freedom (SSF) Data Management System (DMS), Communication and Tracking System (C&T), Power System, etc. The 386DX is also a widely used processor in today's personal computer market.

The 386DX is a 32-bit microprocessor designed to support a multitasking operating system such as UNIX. There are eight general purpose 32-bit registers with 8-, 16-, or 32-bit data types. The 386DX addresses 4 gigabytes (2^{32}) of memory and has two modes of operation: real address mode (real mode) and protected virtual address mode (protected mode) [5].

The 386DX has a on-chip memory management unit (MMU) that supports virtual memory management, and four levels of protection for isolating and protecting applications and the operating system from each other. In addition, the 386DX allows the simultaneous running of multiple operating systems.

The 386DX supports pipelined instruction execution, and the commercial chips are available at clock rates of 20, 25, and 33 MHz. The corresponding internal bus bandwidths are 40, 50, and 66 megabytes per second. The 386DX is upward compatible with the Intel 80X86 family of microprocessors.

A floating point coprocessor, the 387DX, designed to work with the 386DX is available. The 387DX expands the 386DX data types to include 32-, 64-, and 80-bit floating point, and 32-, 64-bit integers.

The 386DX and 387DX were qualified for the MIL-STD-883C Class B qualification (for military applications) in 1989. Intel plans to have the 25 MHz 386DX and 387DX meet the Class S qualification (for space applications) in the third quarter of 1990 [6].

The 386SX and 387SX have been selected to be used in the Multiplexer-DeMultiplexer (MDM) subsystem in the current design of the SSF DMS, instead of the Intel 80186. The primary reason is that the 386SX can run the existing large 386DX software in a cost effective 16-bit hardware environment.

The 386SX is very similar to the 386DX except: (1) the 386SX data bus is 16-bit; (2) the 386SX clock rate is 20 MHz; (3) the bus bandwidth is 20 megabytes per second; (4) the 386SX

addresses up to 16 megabytes (2^{24}) of memory; (5) the matching floating point coprocessor for the 386SX is the 387SX; (6) the 386SX power consumption is lower than the 386DX.

The MIL-STD-883C qualification schedules for the 386SX and 387SX are the same as the 386DX and 387DX.

The power dissipation data for the above processors are listed in Table 1 (the data are estimated from power supply current assuming 5 V nominal voltage).

Microprocessor	Power Dissipation
386DX (20, 25, 33 MHz)	2.5, 2.8, 2.8 W
387DX (20, 25, 33 MHz)	1.5, 2.0, 2.0 W
386SX (20 MHz)	1.5 W
387SX (20 MHz)	1.5 W

Table 1. Power dissipation of the 386/387DX and 386/387SX processors

3-2. Intel i486 Microprocessor

The i486 is currently the top-of-the-line processor in the Intel 80X86 family. The i486 is binary compatible with the 386DX and 387DX processors. The i486 provides a performance level similar to today's RISC processors, yet retains compatibility with a large base of existing software.

The on-chip integration of the i486 includes a floating point unit (FPU), a paged, virtual memory management unit (MMU), and an 8 Kbyte code and data cache. The on-chip cache is fundamental to achieving the high performance level of the i486. Though the 8-Kbyte cache is considerably smaller than the 32-Kbyte external caches built into many 386-based PCs, it is considerably more sophisticated. Wide buses from the cache let the instruction prefetcher grab blocks of code 16 bytes at a time, instead of four. In contrast to most processors with on-chip caches, the i486 cache is unified to hold both code and data. This generally provides more efficient cache utilization than would, for example, separate 4-Kbyte code and data caches. Intel claims that it solves, quite cleanly, problems that might otherwise arise in executing application programs that include self-modifying code [5].

The disadvantage of a unified cache design is that code and data fetches can collide, stalling instruction execution. The i486 prefetch system avoids this hazard in two ways. Instructions are

fetched 16 bytes at a time which reduces the frequency of collisions, and when collisions do occur, the data request is serviced first [7].

The i486's on-chip cache uses a write-through memory update policy (see Section 2-2 for the characteristics of this policy). Because of the small cache size and write-through memory policy, shared-memory multiprocessor systems may require second-level caches.

The i486 is a new microprocessor from Intel. A few errors were found in this complex 1.2-million transistor chip. The quality and reliability concerns may be an issue at this time. Intel plans to have the 25 MHz i486 meet the MIL-STD-883C Class B qualification (for military applications) in 1992, and Class S qualification (for space applications) in 1993 [6].

The power dissipation data for the i486 are listed in Table 2 (the data are estimated from power supply current assuming 5 V nominal voltage).

Microprocessor	Power Dissipation
i486 (25, 33 MHz)	3.5, 4.5 W

Table 2. Power dissipation of the i486 processor

3-3. Intel i860 Microprocessor

The Intel i860 microprocessor is the first RISC-based chip from Intel. Intel claims that the i860 is a 64-bit microprocessor because the external data bus, the internal instruction cache bus and the 3-D graphics unit are all 64-bit. The internal data cache bus is 128-bit [5].

The overall level of integration in the i860 is quite remarkable. It incorporates on one chip: (1) a 32-bit integer and control unit; (2) pipelined floating-point control, adder, and multiplier units; (3) a 64-bit 3-D graphics unit; (4) a 4 Kbyte instruction cache; (5) an 8 Kbyte data cache; (6) a paged, virtual memory management unit; and (7) a bus control unit. The i860 has 32 integer registers and 32 floating-point registers, each 32 bits wide. The floating-point registers are also used by a set of graphics operations, primarily for 3-D graphics computations.

The i860 can execute up to three instructions at the same time per clock cycle: one integer or control instruction per cycle, and up to two floating-point results per cycle. The i860 has 33 and 40 MHz clock speeds available today, 80 MHz in the future. The i860 performance is most impressive on vectorizable floating-point programs and 3-D graphics calculations. As an integer machine, it does not have any architectural advantage over other RISC processors.

For engineering, scientific, robotic and automation applications, such as modeling a robotic arm movement, data must be presented so that the user's innate ability to process visual information is tapped. This means generating a realistic, 3-D image rather than columns of floating-point numbers. To facilitate data visualization, the i860 instruction set includes graphics primitives, such as the z-buffer, aimed at increasing the throughput of 3-D graphics applications.

Though Intel claims the i860 as a stand-alone, high-performance architecture, there are some factors that limit the i860's present use as a general-purpose CPU. These factors are the currently limited available software and the incomplete instruction set . For example, control instructions are only able to test a single status code bit, and there is no add-with-carry or subtract-with-borrow instruction, etc.

The i860 will probably be most popular as a graphics or floating-point coprocessor in systems based on another processor. For example, the DECstation 5000 has a MIPS R3000 as the CPU, and an i860 as a graphics coprocessor.

Used as a coprocessor for the 386DX/i486, the i860 may be useful for the SSF DMS, or other space-related projects. The i860 can complement the weakness of the 386DX/i486 in graphics applications. The capability of 3-D solid modeling and movement are also essential for efficient robotic and automation applications.

Intel plans to have the i860 meet the MIL-STD-883C Class B qualification in the second quarter of 1991 [6].

The power dissipation data for the i860 are listed in Table 3 (the data are estimated from power supply current assuming 5 V nominal voltage).

Microprocessor	Power Dissipation
i860 (33, 40 MHz)	3.0, 3.3 W

Table 3. Power dissipation of the i860 processor

3-4. Intel 80960 Embedded Processor

The Intel 80960 family of embedded processors is a new 32-bit architecture from Intel. The family includes four processors: the 960KA, 960KB, 960MC, and 960CA. The on-chip integration and clock rate for each processor are showed in Table 4 [8].

	960KA	960KB	960MC	960CA
on-chip floating-point unit	No	Yes	Yes	No
on-chip memory mngt. unit	No	No	Yes	No
on-chip cache (byte)	512	512	512	1K instruction
on-chip RAM (byte)	0	0	0	1K data
on-chip DMA channel	0	0	0	4
clock rate (MHz)	25	25	20	16,25,33

Table 4. On-chip devices and clock rates of the 80960 family processors

The key feature of the 960 family is the superscalar implementation (IBM's new RS/6000 uses many of the same techniques). On each clock cycle the 80960 prefetches either three or four instruction words from the on-chip cache, depending on whether the current instruction pointer is odd or even.

Intel's marketing literature and technical presentations claim that the 960 family is the first microprocessor that can sustain execution rates "up to two instructions per clock (IPC) within optimized program loops". Actually, the 960 has difficulty exceeding 2 IPC due to the bus interface bottleneck. The overall throughput of a real 80960-based system is likely to vary widely according to the nature of the application [3].

The 960 processors are intended for use in embedded applications demanding very high computational power. These include robotics, avionics, machine control, process control, laser printers and image processing equipment, etc. The 960CA has four on-chip Direct Memory Access (DMA) channels to make it suitable for high speed LAN (e.g. FDDI) communications controllers.

The 960MC processor is designed for military applications and can operate in temperatures from -55°C to +125°C. The 960MC is a strong candidate to replace the 16-bit MIL-STD-1750A instruction set architecture. (Please see the section of the MIL-STD-1750A architecture and RH-32

program for details). Intel plans to have the 960MC meet the MIL-STD-883C Class B qualification in September, 1990 [6].

The power dissipation data for the 80960 family are listed in Table 5 (the data are estimated from power supply current assuming 5 V nominal voltage).

Microprocessor	Power Dissipation
960KA (16, 20, 25 MHz)	1.9, 2.1, 2.5 W
960KB (16, 20, 25 MHz)	1.9, 2.1, 2.5 W
960MC (20 MHz)	2.8 W
960CA (16, 25, 33 MHz)	1.9, 2.8, 3.5 W

Table 5. Power dissipation of the 80960 family processors

4. MIPS PROCESSORS

4-1. MIPS R2000 and R3000 Microprocessors

The MIPS processor design evolved from a project called MIPS (Microprocessor without Interlocked Pipeline Stages) at Stanford University in the early 1980's under the direction of Prof. John Hennessy.

The MIPS R2000 and R3000 processors share an identical instruction set architecture. The two processors are 32-bit RISC-based microprocessors with 32 general purpose 32-bit registers. All instructions and addresses are 32-bit. The R2000/R3000 processor has an on-chip memory management unit and the address space is 4 Gbyte [9].

The R3000 has a higher clock rate implementation than the R2000: the clock rate for the R2000 is 16 MHz, while the R3000 has 16, 20, and 25 MHz rates. Both processors have on-chip separate instruction and data caches, but the sizes are different: the R2000 has 4 to 64 Kbytes instruction and data caches each, while the R3000 has up to 256 Kbytes each [10].

The R2000/R3000 architecture defines both integer and floating-point instruction and register sets. Floating-point support is provided by a tightly-coupled coprocessor. The floating-point coprocessor for the R2000 is the R2010, for the R3000, the R3010. The R2000/R3000 is especially strong in floating-point performance. MIPS Computer claims the 25-MHz R3000/R3010 pair is rated at 17.3 MWhetstones single-precision and 13.6 MWhetstones double-precision.

Several RISC processors, including the R2000/R3000, the MC88000, etc., support both big-endian and little-endian byte ordering schemes (please refer back to Section 2.3 for definition). Ability to match the byte ordering of existing systems is important for retaining data compatibility, and eases porting software from earlier architectures. The R2000/R3000 can operate in either mode, but this is determined by the state of a pin at reset, and can not be dynamically changed while the processor is running. This problem is solved in the MIPS II architecture and the R6000 processor, which is discussed in the next section.

The MIPS architecture does not include an indivisible test-and-set or swap instruction which is useful for multiprocessor support. Other RISC-based architectures, such as i860, MC88000, and SPARC, all have some type of multiprocessor support.

The R2000/R3000 processor uses CMOS (Complementary metal-oxide semiconductor) technology. MIPS Computer contracted NEC of Japan, Siemens, LSI Logic, Performance Semiconductor, and Integrated Device Technology (IDT) to manufacture chips.

The R2000/R3000 processor is qualified for military temperature (0°C to +70°C) and MIL-STD-883C Class B devices. The R3000-based instruction set architecture (ISA) is a strong candidate considered to replace the traditional 16-bit MIL-STD-1750A standard. (Please see the section of the MIL-STD-1750A architecture and RH-32 program for details).

The power dissipation data for the above processors manufactured by the LSI Logic are listed in Table 6.

Microprocessor	Power Dissipation
R2000 (16 MHz)	3 W
R2010 (16 MHz)	3 W
R3000 (16, 20, 25 MHz)	3, 3.5, 4 W
R3010 (16, 20, 25 MHz)	3, 3.5, 4 W

Table 6. Power dissipation of the R2000/R2010 and R3000/R3010 processors

4-2. MIPS R6000 Microprocessor

The R6000 is a RISC-based 32-bit microprocessor. The R6000 uses ECL (Emitter-coupled logic) technology. Generally speaking, ECL is used in systems requiring high-speed operations, and CMOS is used in systems requiring low power consumption, which is almost essential for embedded space applications. The much higher power dissipation data of the R6000 are listed at the end of this section. MIPS Computer contracted with NEC and Sony to manufacture the R6000.

The R6000 supports the MIPS II instruction set architecture, which includes load interlocks and a number of new instructions [11]. This instruction set is a superset of the MIP I ISA, which is implemented by the R2000 and R3000 processors. The R6000 is upward object-code compatible with existing MIPS processors.

The R6000 has the capability of switching between big-endian and little-endian modes dynamically, i.e., while the processor is running. (The R2000/R3000 can also operate in either mode, but this is determined statically before a processor is started.)

The matching floating-point coprocessor for the R6000 is the R6010. The available clock speed is 66.7 MHz, up to 80 MHz planned in 1990. MIPS Computer claims the R6000 has 55 VAX MIPS and 32 MWhetstones (double-precision) performance.

The power dissipation data of the R6000/R6010, which are much higher than the R2000/R3000 due to the ECL process, are listed in Table 7.

Microprocessor	Power Dissipation
R6000 (66.7 MHz)	16 W
R6010 (66.7 MHz)	16 W

Table 7. Power dissipation of the R6000 processors

5. SUN ARCHITECTURE

5-1. Sun SPARC Architecture

Sun Microsystems' SPARC (Scalable Processor ARChitecture) design is derived from the RISC work done at the University of California at Berkeley from 1980 through 1982 under the direction of Prof. David Patterson. The SPARC architecture, Version 1, was introduced in March, 1986. The current version is Version 7.

As an architecture, SPARC is not a particular chip or implementation. The SPARC architecture defines a set of 32-bit instructions, a set of registers, how the registers work, and how traps and interrupts work. The SPARC architecture does not define details such as the size and timing of data and address busses, caches, or memory management units [12].

The SPARC processor logically comprises an integer unit and a floating-point unit. The SPARC follows the Berkeley RISC design philosophy in that it stresses the importance of the CPU register file and implements similar register window features. The number of windows that can be implemented on different versions of the SPARC ranges from 2 to 32, resulting in a total number of integer unit general-purpose registers from 40 to 540, respectively. The purpose of the large, windowed register file is to allow compilers to cache local values across subroutine calls and provide a register-based parameter passing mechanism. Since the SPARC architecture is a register-intensive architecture, most instructions operate on either two registers or one register and a constant, and place the result in a third register. Only load and store instructions access storage.

The SPARC architecture supports tagged arithmetic instructions, which is unique among the RISC-based processors. The tagged instructions can be used by languages such as Lisp, Smalltalk, and Prolog that benefit from tags. The tagged instructions provide the capability to tag data and pointers differently so that detection of illegal operations on the data or pointers can be performed during execution.

The SPARC architecture supports two types of semaphore instructions (though early implementations only support one). The Load-Store Unsigned Byte instruction reads a memory location and then writes that memory location to all 1s in an atomic manner. The Swap instruction causes a memory location to be read and then replaced with the contents of a specified register.

The SPARC architecture supports the big-endian byte-ordering format only. The SPARC originated at Sun Microsystems Inc., where most of the products are Motorola MC680X0-based

(big-endian byte order). Big endian format thus became the logical choice. Most other RISC-based processors, such as MIPS, i860, and MC88000 support big-endian and little-endian byte ordering formats.

The SPARC instruction set does not provide basic integer multiply and divide instructions. Alternatively, it provides multiply-step and divide-step instructions, and library routines to implement both multiply and divide operations. As such, the SPARC performance will suffer on applications that require extensive multiplication and division operations unless vendors add the basic multiply and divide instructions to the architecture. The lack of these instructions limits the SPARC architecture ability to increase multiplication and division performance when using hardware available in future implementations.

Sun was concerned with getting their first SPARC chip into production as rapidly as possible, and thus designed the processor in a gate-array technology. Sun licensed many semiconductor companies to manufacture chips, such as Systems & Processes Engineering Corporation (SPEC), Cypress, LSI Logic, Fujitsu, Bipolar Integrated Technology (BIT), etc.

SPEC has been funded by a NASA SBIR (Small Business Innovation Research) Program to develop a gallium arsenide (GaAs) SPARC processor to demonstrate the inherent speed and radiation-hardness advantages of GaAs. Multiple GaAs SPARC processors will be included in a demonstration board that SPEC is building to look at the GaAs capabilities. The board will include four to eight GaAs SPARC processors, GaAs array communications coprocessors and GaAs floating point coprocessors. SPEC claims that the GaAs SPARC processors and coprocessors are being designed to operate at 200 MHz, with performance at 800 to 1600 MIPS in a four- to eight-processor implementation. Initial samples of the GaAs SPARC processor is scheduled late in 1990.

Cypress has implemented a SPARC chip (called CY7C601) qualified for the MIL-STD-883C Class B (for military applications) at 33 MHz. If requested, Cypress will contract TRW to manufacture a Class S (for space applications) SPARC chip.

The power dissipation data of the Cypress SPARC chip, CY7C601, is listed in Table 8.

Microprocessor	Power Dissipation
Cypress CY7C601 (33 MHz)	3 W

Table 8. Power dissipation of the Cypress CY7C601 SPARC processor

6. MOTOROLA PROCESSORS

6-1. Motorola MC68030 Microprocessor

The MC68030 is a 32-bit, CISC-based microprocessor from Motorola. The MC68030 has been widely used in desktop workstations and computers such as Apple's Macintosh computers, NeXT computer, and many Hewlett-Packard workstations.

The MC68030 on-chip devices include a central processing unit (CPU), 256-byte instruction and data caches each that can be accessed simultaneously, a bus controller that supports synchronous and asynchronous bus cycles and burst data transfers, and a memory management unit (MMU) [13].

The MC68030 on-chip MMU, which allows page sizes from 256 bytes to 32 Kbytes and supports a variety of page table formats, evolved from the external MMU chip for the MC68020. This variety of page sizes increases the flexibility of the MMU. (This flexibility has been reduced in the MC68040 MMU, please see the MC68040 section for comparison).

The MC68030 bus supports 8-, 16-, and 32-bit memories and peripherals to make it upward object code compatible with the earlier members of the MC68000 family. The MC68030 has 4-Gbyte addressing range. The available clock rates are 20, 25, 33, and 40 MHz.

The MC68030 uses a three-stage pipelined internal architecture to improve throughput. The pipeline allows as many as three words of a single instruction or three consecutive instructions to be decoded concurrently.

Motorola has the MC68030 qualified for the MIL-STD-883C Level B (military applications), but Motorola does not plan to manufacture a Class S (space applications) qualified chip at this time [14].

The power dissipation data of the MC68030 is listed in Table 9.

Microprocessor	Power Dissipation
MC68030 (33 MHz)	3 W

Table 9. Power dissipation of the MC68030 processor

6-2. Motorola MC68040 Microprocessor

The MC68040 is Motorola's top-of-the-line microprocessor in the 32-bit MC68000 family of processors. Motorola claims the MC68040 is 100% compatible with existing MC68000 family processors and operates at about three times the speed of the MC68030.

The MC68040 includes an integer unit, a floating-point unit, independent 4 Kbyte instruction and data caches, and dual independent memory management units (MMUs) [15].

The MC68040 memory management system is sophisticated. As has become standard, the MC68040 supports a demand-paged, memory management environment. Unlike most other processors with the exception of the MC88000, the MC68040 has *separate* MMUs so that simultaneous data and instruction accesses do not contend for a single MMU. The elimination of contention increases the potential for parallelism.

One area where the flexibility of the MMU was reduced (see the MC68030 section for comparison) was the page size: only 4K-byte and 8K-byte pages are supported by the MC68040 MMU. The page size must be reasonably large to allow the overlap of cache access and address translation. The page size could have been as small as 1K bytes, but the trend in operating systems is toward larger pages.

A high degree of instruction execution parallelism is achieved through the use of multiple independent execution pipelines, multiple internal buses, and the separate physical caches for both instruction and data accesses. Cache functionality is enhanced by the inclusion of on-chip bus snooping logic to support cache coherency.

The MC68040's caches and bus support a sophisticated snooping protocol, but it does not support a full multiprocessor protocol, which requires that the caches be able to identify data as shared or exclusive.

The power dissipation data of the MC68040 is listed in Table 10.

Microprocessor	Power Dissipation
MC68040 (25 MHz)	5 W

Table 10. Power dissipation of the 68040 processor

6-3. Motorola MC88000 Family of Microprocessors

The MC88000 family of processors is the first RISC-based processors introduced by Motorola. The MC88000 chip set includes the MC88100 CPU and the MC88200 Cache/Memory Management Unit (CMMU).

The MC88100 includes an integer unit, a floating-point unit, a register file with thirty-two 32-bit, general-purpose registers, an instruction unit, and a data memory unit. All instructions are implemented directly in hardware [16].

The MC88100 has four independent execution units (integer, floating-point, execution, and data units) with five concurrent pipelines. All these execution units contain an additional level of fine-grain parallelism. The three internal register buses allow three simultaneous register accesses, eliminating internal bus contention by the concurrent execution pipelines. The MC88100 also supports concurrent register writes and reads. One execution unit can write a result to the destination register while another unit fetches the source operands from other registers. The pipelined load and store operations can reach up to 80 Mbytes per second at 20 MHz.

While the MC88100 memory ports can interface directly to memory, most MC88100 designs incorporate at least two MC88200s, one for data memory and one for instruction memory. The data unit and instruction unit use the separate processor buses to interface these units to the respective MC88200s/memories. Two to eight CMMUs can be incorporated into an MC88100 system [17].

The MC88200 has 16 Kbytes of cache memory and is designed to provide the mapping and protection needed to construct a multitasking, demand-paged memory system. The MC88200 resides logically between the processor and the physical memory. In this configuration, the MC88200 controls all accesses to physical devices, and tasks can be prevented from accessing the resources owned by other tasks.

Motorola plans to have the MC88000 family qualified for the MIL-STD-883C Level B (military applications) in the third quarter of 1990, but Motorola does not plan to have the chip qualified for Level S (space applications) at this time [14].

The power dissipation data for the MC88000 are listed in Table 11.

Microprocessor	Power Dissipation
MC88100 (25 MHz)	1.5 W
MC88200 (25 MHz)	1.5 W

Table 11. Power dissipation of the MC88100 and MC88200 processors

7. IBM PROCESSOR

7-1. IBM RS/6000 Microprocessor

IBM introduced a series of new RISC-based workstations and file servers in February, 1990. The architecture is called POWER (Performance Optimized With Enhanced RISC) and the product line is called RISC System/6000 (RS/6000) by IBM.

The RS/6000 processor architecture is based on a logical view of a processor, which consists of three independent functional units: an integer unit, a floating-point unit, and an instruction cache unit. IBM has taken a unique approach to processor partitioning. The integer and floating-point units are each separate chips, which is not out of the ordinary. What is unusual is that instruction decoding and branch processing logic is not part of the integer processing chip, but is rather part of the instruction cache chip [18].

The RS/6000 has superscalar implementation. Four instructions at a time are fetched from the cache, and under the right circumstances, all four can be dispatched in a single cycle. Up to five operations can be executed per cycle: one branch, one fixed-point, and three floating-point instructions.

The RS/6000 has a 4 Gbyte address space. The page size is 4 Kbytes. The data and instruction caches are separate. There are three clock speeds available for the RS/6000 series: 20, 25, and 30 MHz. IBM claims the performance benchmarks of 27.5, 34.5, and 41.1 MIPS (based on Dhrystone 1.1), and 7.4, 10.9, and 13 MFLOPS (based on double-precision Linpack, Fortran compiled), respectively. The Micro Channel with enhanced streaming mode, as implemented on the RS/6000, provides 20 to 25 Mbytes per second of I/O bandwidth.

The RS/6000 is a new product. How successful the product will be in the commercial market and the availability of a space qualified chip are unknown at this time.

8. JAPAN'S TRON PROJECT AND PROCESSOR

8-1. Japan's TRON Project

TRON is a major project now under way in the Japanese microcomputer world. Although its name is an acronym for The Real-time Operating System Nucleus, TRON's goal is not to create simply a real-time operating system. Rather, it is to create an integrated, open design for all aspects of microcomputer architecture and use - from microprocessor to bus, LAN, operating system, user interface, and application languages. Because NASDA (National Space Development Agency of Japan) plans to use TRON-based microprocessor and operating system in the Space Station Freedom (SSF) Japanese Experiment Module (JEM), it is logical to discuss the TRON project in this report.

Dr. Ken Sakamura initiated the TRON project in 1984. TRON's fundamental goal is to create the architecture for a computer environment that both technologically and legally achieves wide acceptance in and out of Japan. The architecture begins with a 32- and 64-bit microprocessor and goes on to define: (1) the real-time operating system for industrial applications, ITRON; (2) a real-time operating system for personal computers and workstations, BTRON; (3) an operating system for central systems, CTRON; and (4) a network operating system to tie these pieces together, MTRON [19].

These elements are designed with the idea that an extremely high level of distributed processing will characterize the future. This highly distributed system is called HFDS (Highly Functional Distributed System). The attempt is to standardize the user interface and to create protocols for interchange and translation of mixed numeric, textual, graphics, motion, and voice data across such networks.

In the TRON project, the interface between the layers of the hierarchical architecture are standardized, but the TRON specifications do not spell out the method of implementation. As long as the interface standard is followed, different layers of different origins can be used together. The TRON project allows for variations in implementation to incorporate advanced technology, and to allow the system to be adapted to various applications while assuring healthy competition between commercial implementors. The TRON project standards are referred to as "weak standards". As this term implies, the TRON architecture is a general architecture and is not tied to any particular hardware or software [20].

8-2. TRON VLSI CPU Architecture

The TRON project intends to build the HFDS with a general-purpose microcomputer chip family and then use this in communication machines and various other objects. In order to have a single underlying CPU, a subproject was begun to design a CPU, which resulted in the specification for the TRON VLSI CPU.

The chip architecture specifications for TRON have been designed from scratch and in parallel with the design of the TRON specifications for operating systems as part of the TRON total architecture. The TRON VLSI CPU has instructions to accelerate certain operations of the ITRON and BTRON operating systems.

The specification uses 32-bit addressing at present, but extension to 64-bit is already defined. The TRON specification provides two classes of instruction formats: a general format and a short format. The length of the short format is two byte, while that of the general format, excluding the addressing extension portion, is four byte. The length of frequently used instructions is short. The big-endian byte ordering scheme is employed in the TRON project [21].

The TRON project standards are referred to as "weak standards", as mentioned in the previous section. The instruction set processor specification for the TRON VLSI CPU specifies the type of instructions, but not the hardware implementation. In early 1988, the first commercial implementation of the TRON VLSI CPU was announced. This was the GMICRO/200 by Hitachi Ltd., which supports virtual memory. There are currently eight microprocessors from six companies based on the TRON specification. Most of these microprocessors have certain amount of on-chip cache memory. The clock speeds range from 20 to 33 MHz.

9. MIL-STD-1750A ARCHITECTURE AND RH-32 PROGRAM

9-1. MIL-STD-1750A Instruction Set Architecture

MIL-STD-1750A was developed by the United States Department of Defense in 1979 to establish a uniform instruction set architecture for airborne computers to be used in Air Force avionic weapon systems. The standard was adopted by the Air Force in 1980. Benefits such as the reduction in total support software required has led to the standard being adopted into many Navy and Army programs.

This standard defines a 16-bit computer instruction set architecture (ISA) based on the DEC PDP-11 ISA. The ISA supports 16-bit fixed point single precision, 32-bit fixed point double precision, 32-bit floating point and 48-bit floating point extended precision data in 2's complement representation [22].

This standard only defines the ISA. It does not define specific implementation details of a computer. A typical implementation of the 1750A includes a general purpose, single chip, 16-bit CMOS microprocessor. The data type could be 16- and 32-bit integer, 32-, and 48-bit floating point numbers.

The implementation is normally designed for extensive error management, fault management, and interrupt capability. Operating from a single 5 V power supply, it consumes less than 500 milliwatts of power over the full military temperature range. The 1750A is also required to meet strict radiation-hardened requirements.

The 1750A ISA was developed more than ten years ago. The standard was popular but has the following disadvantages comparing to modern processors: (1) The word length of the 1750A ISA is only 16-bit. It has less capability than the modern 32-bit processors. (2) The addressing capability of the 1750A is only 16 Megabytes (2^{24} bytes), which is much less than the 4 Gigabytes (2^{32} bytes) capabilities of the modern processors. (3) There is much less available commercial-off-the-shelf software for the 1750A than the modern computers.

9-2. RH-32 Program

The U. S. Air Force has realized the limitations of the MIL-STD-1750A ISA and is looking at options to replace the 1750A. The Radiation-Hardened 32-Bit Processor (RH-32) program was created in May, 1988 to build a rad-hard, MIPS R3000 ISA-based or Intel 80960MC ISA-based processor.

The requirements for the RH-32 processor include: (1) support spaceborne and airborne avionics real-time fault tolerant processing; (2) efficiently execute Ada software; (3) 20 MIPS (core ISA) sustained throughput in a worst case, post-radiation environment; (4) directly address 4 gigabytes (2^{32} bytes) of memory; (5) 32-bit wide data paths (nominally); (6) IEEE-STD-754 binary floating point arithmetic; and (7) standard interfaces.

Phase I of the RH-32 program was scheduled from May, 1988 to June, 1989. Four contractors were selected for the Phase I project: TRW, Honeywell, Unisys, and IBM. The products of Phase I include: (1) standard evaluation circuits; (2) microprocessor hardware and software designs, which include preliminary microcircuit specifications, operational concept, software support tool specifications, and interface to Ada environment; (3) instruction set architecture models, simulators, and benchmarks.

Phase II was scheduled from July, 1989 to June, 1991. The Air Force and program members evaluated the deliverables from the four Phase I contractors and selected TRW and Honeywell for the Phase II project. The products of Phase II include: (1) rad-hard microprocessor chips; (2) advanced development models (ADM), which includes nine in spaceborne configuration and one in avionics configuration; (3) software tools, which includes support tools (core ISA to machine executable code) and interface tools (interface ADMs to host computer); (4) MIL-STD-883C Class S (space application) qualification plan.

10. COMPARISONS

10-1. Comparisons of Intel 386/387DX and i486 Microprocessors

A. OBJECTIVE

As mentioned before, the 386/387DX have been baselined for the SSF DMS, but the i486 has superior performance to the 386/387DX. This section compares their ISA, performance, power consumption, and MIL-STD-883C Class S qualification schedule.

B. COMPARISON OF INSTRUCTION SET ARCHITECTURE

The i486 implements essentially the same instruction set, programming model, memory organization, process organization, process management facilities, and so forth, as the original 386/387DX, though the i486 significantly outperforms the 386/387DX. Intel made no significant changes to the 386/387DX ISA except for the addition of three new instructions required for cache support (INVLPG, INVD, and WBINVD) and three for multiprocessing functions (CMPXCHG, XADD, and BSWAP).

C. COMPARISON OF PERFORMANCE

Intel claims that performance of the i486 at 25 MHz will be about 37,000 Dhrystones, 6.1 M Whetstones (double precision), or about 15-20 VAX MIPS. Intel also claims a performance enhancement of 2 to 4 times over the 386/387DX. The fully redesigned integer unit and on-chip cache reduce the average number of clocks per instruction by a factor of about 2.5. The on-chip FPU executes floating-point instructions about three times faster than the 386/387DX combination. And while the initial i486 clock speed is 25-MHz, the part was designed to scale to clock rates higher than those achievable with the 386/387DX. The comparisons of instruction execution time between the 386/387DX and i486 are listed in Table 12 [7].

Instruction	Optimum Clock Cycle Count	
	386/387DX	i486
Register-to-register add	2	1
Memory load (16 bit)	2	1
Memory store (16 bit)	2	1
Memory-to-register add (16 bit)	6	2
Register-to-memory add (16 bit)	7	3
Unconditional jump	8	3
Call	8	3
Return	11	5

Table 12. Comparisons of the 386/387DX and i486 instruction execution time

D. COMPARISON OF POWER CONSUMPTION

As mentioned before, the i486 contains both integer unit (compatible with the 386DX) and floating-point unit (compatible with the 387DX). The 25-MHz i486 power dissipation (3.5W, as listed in the Table 13) is 1.3W lower than the sum of the 386DX and 387DX (4.8W), and the i486's 3.5W even includes an on-chip 8-Kbyte cache. Therefore, the strict power consumption requirement for the SSF DMS may not cause a problem for using the i486, instead of the 386/387DX. The power dissipation data of the 386/387DX and i486 processors are listed in Table 13.

Microprocessor	Power Dissipation
386DX (25, 33 MHz)	2.8, 2.8 W*
387DX (25, 33 MHz)	2.0, 2.0 W*
i486 (25, 33 MHz)	3.5, 4.5W*

Table 13. Power dissipation of the 386DX, 387DX, and i486 processors

* Estimated from power supply current of the Intel *Microprocessors* data book, 1990. These figures are different from the data book published by Intel in 1989.

E. CLASS S QUALIFICATION SCHEDULE

As mentioned before, Intel plans to have the 25 MHz 386DX and 387DX meet the MIL-STD-883C Class S qualification (for space applications) in the third quarter of 1990. Intel also plans to have the 25 MHz i486 meet the Class B qualification (for military applications) in 1992, and Class S qualification at the end of 1993.

10-2. Comparisons of the i860, R3000, SPARC, and MC88000

The four RISC-based processors are compared and listed in table 14. ("A" is slightly superior to the others, "B" is essentially equal to the others, and "C" is slightly inferior with respect to the others. The data for the IBM RS/6000 is insufficient for comparisons.)

	i860	R3000	SPARC	MC88000
Registers	B	C	A	C
Multiply/divide instruction	B	A	C	A
Flexibility of byte ordering	B	A	C	B
Memory management	C	B	A	B
Multiprocessing support	A	C	A	A
Supporting software	C	B	A	C
Graphics applications	A	B	B	B
AI applications	B	B	A	B
MIL-STD-883C Class B	B	A	A	A
MIL-STD-883C Class S	B	B	B	C

Table 14. Comparisons of the i860, MC88000, R3000, and SPARC

Registers: The SPARC supports a register window system. The i860 has 32 integer and 32 floating-point registers. The MC88000 and R3000 have 32 registers for both integer and floating-point operations.

Multiply/divide instruction: The MC88000 and R3000 provide both of the basic integer and divide instructions. The i860 supplies a multiply operation via its floating-point operation, but provides a library routine for division. The SPARC provides a multiply step instruction and library routines to implement both multiply and divide operations.

Flexibility of byte ordering: The MIPS II in the R6000 supports both of the big endian and little endian byte ordering schemes dynamically. The i860 and MC88000 support both byte ordering schemes statically. The SPARC supports the big endian scheme.

Memory management: All of the architectures support a 4-Kbyte page size. The SPARC supports 256-Kbyte, 16-Mbyte, and 4-Gbyte mappings. The R3000 provides 256-Kbyte mappings. The MC88000 supports 4-Mbyte mappings. The i860 does not support any form of larger mappings.

Multiprocessing support: The i860, MC88000 and SPARC support some kind of semaphore or atomic test-and-set type of instruction. The R3000 does not have an atomic test-and-set or exchange instruction for multiprocessing support.

Supporting software: The SPARC has the largest commercial-off-the-shelf supporting software. The R3000 has some supporting software. The i860 and MC88000 have little supporting software available for these new architectures.

Graphics applications: The i860 has an on-chip 3-D graphics processor and z-buffer implementation which makes it superior to other three processors in graphics applications.

Artificial intelligence applications: The SPARC supports tagged arithmetic that makes it superior to the other three processors in artificial intelligence applications.

MIL-STD-883C Class B and S: Intel plans to qualify the i860 for the Class B and S between 1991 and 1993. The R3000, Cypress SPARC and MC88000 are qualified for Class B. The MC88000 has less possibility to be space-qualified.

11. FUTURE TRENDS OF MICROPROCESSORS

From the above discussions, the following future trends of microprocessors are predicted:

(1) Many more transistors on a chip: The transistor size has been shrinking almost every year. More and more transistors will be able to be installed in a microprocessor. For example, Intel i486 has about 1.2 million transistors on a chip. Intel predicted that a 100-million transistor chip may be available by the year 2000 [23].

(2) Much higher clock rate: The process technology of microprocessors has changed from PMOS, NMOS, CHMOS, to ECL, etc. The clock frequency has exceeded 66 MHz in some implementations. We can expect even higher clock frequency as technology grows. Intel predicted that a 250 MHz clock chip may be available by the year 2000.

(3) Many more instructions per cycle: More and more instructions will be able to be processed in a single clock cycle. The combination of higher clock rate and more instructions per cycle execution can increase the processor performance.

(4) Multiple CPUs in a single chip: The semiconductor industry has been able to incorporate a central processor and a floating-point coprocessor in a single chip. The next generation of microprocessors may have four, eight, or even more CPUs in a chip. This may have a significant impact on parallel processing technology.

(5) Large on-chip cache: The amount of on-chip cache has been increasing from 256 bytes, to 8 Kbytes, and even 256 KBytes each for instruction and data caches (i.e. the MIPS R3000). The on-chip cache may reach several Mbytes in the future.

12. REFERENCES

1. Piepho, S. Richard; and Wu, S. William: A Comparison of RISC Architectures, IEEE Micro, August 1989, pp. 51-62.
2. Almasi, S. George; and Gottlieb, Allan: Highly Parallel Computing, Benjamin/Cummings, 1989, pp. 414-419.
3. Wharton, H. John: Conflict Avoidance Key to Superscalar Efficiency, Microprocessor Report, November 1989, pp. 22-25.
4. U.S. Air Force: Military Standard Test Methods and Procedures for Microelectronics, MIL-STD-883C (USAF), 12 February, 1988.
5. Intel Corporation: Microprocessors, 1990.
6. Gregory L. Mather, Intel military/space programs manager, private communication.
7. Wharton, H. John: The Intel 80486 Strikes Back, Microprocessor Report, April 1989, pp. 8-12.
8. Intel Corporation: 16-/32-Bit Embedded Processors, 1990.
9. Kane, Gerry: MIPS R2000 RISC Architecture, Prentice Hall, 1987, pp. 2-1 to 2-15.
10. MIPS Computer Systems, Inc.: MIPS R3000 Processor Interface, 1989.
11. MIPS Computer Systems, Inc.: MIPS R6000 Interface Specification, 1989.
12. Sun Microsystems, Inc.: The SPARC Architecture Manual, Version 7, 1989.
13. Motorola, Inc.: MC68030 Enhanced 32-Bit Microprocessor User's Manual, third edition, 1990.
14. Jeffery Mitchell, Motorola technical staff, private communication.
15. Motorola, Inc.: MC68040 32-Bit Third-Generation Microprocessor User's Manual, 1989.
16. Motorola, Inc.: MC88100 RISC Microprocessor User's Manual, second edition, 1990.
17. Motorola, Inc.: MC88200 Cache/Memory Management Unit User's Manual, second edition, 1990.

18. International Business Machines Corporation: RISC System/6000 Technology, 1990, pp. 2-23.

19. Sakamura, Ken: Recent Trends, IEEE Micro, April 1988, pp. 10-11.

20. Sakamura, Ken: The TRON project, Microprocessors and Microsystems, October 1989, pp. 493-502.

21. Sakamura, Ken; and Enomoto, Tatsuya: 32-bit microprocessors based on the TRON specification, Microprocessors and Microsystems, October 1989, pp. 503-513.

22. U.S. Department of Defense: Military Standard Sixteen-Bit Computer Instruction Set Architecture, MIL-STD-1750A (USAF), 2 July, 1980.

23. Intel Corporation: Systems Design in the 90's, 1990.

N93-72608

DMS Processor Evolution Study

Yuan-Kwei Liu

The attached report, written by Y.K. Liu/Code FII of NASA Ames Research Center, is a portion of the DMS EVOLUTION REPORT being prepared by a group of research scientists at Code FII.

PRELIMINARY REPORT

September 1990

Ames Research Center

Moffett Field, California 94035

5. DMS Evolution and Growth

5.2 Directions for Growth

5.2.1. Hardware Component Enhancements

5.2.1.1. Standard Data Processor

5.2.1.1.1 Rationale

The Intel 386DX microprocessor, selected to be the main processing resource for the DMS SDP, has much commercial off-the-shelf (COTS) supporting software. However, the 386DX was proposed by the Work Package II contractors in 1987, and the 4 MIPS performance of the 386DX is relatively low comparing to other processors available today. Further, the 386 is 100% utilized as demonstrated in the DMS End-To-End Simulation Analysis Report (MDC H4875) published by McDonnell Douglas Space Systems Company (MDSSC) in April 1990. This report stated "In fact, the amount of processing required by the application EDP cannot be realized unless its processing rate.....is at least 8 MIPS. However, this still results in relatively large queue sizes and a high utilization." on Page 5-23.

5.2.1.1.2 Objectives

Increasing the processor performance and capability has not only become a wish list item for the DMS, but also a necessity. There are many commercially available processors which have superior performance compared to the 386, but we can not only consider the performance when selecting a processor for the DMS. The processor is the "foundation" of the DMS. It will affect the local bus, system bus, interface unit, global network, etc. that have been selected for the DMS. Besides, once the processor instruction set architecture (ISA) is selected, all of the important system software will be implemented based on this ISA. Selecting an ISA that has a strong commercial support and can be upgraded in the 30-year life cycle of the Space Station Freedom is one of the most important item for the DMS.

PRECEDING PAGE BLANK NOT FILMED

There are many commercial processors that support different ISAs, such as Sun SPARC ISA, MIPS ISA, Motorola 680X0 ISA, 88000 ISA, etc. It is not within the scope of this report to compare these ISAs, but the Intel 80X86 ISA is the most widely used in the United States and has more available supporting software than any other ISA.

5.2.1.1.3 The i486 Features And Advantages

If we decide to stay with the Intel 80X86 ISA, then the i486 seems to be the most reasonable candidate to upgrade the SDP. The i486 has the following important features:

(1) 100% binary compatible with the 386: The i486 ISA contains six more instructions than the 386, and is a super set of the 386 ISA.

(2) Higher performance: The floating point and integer performance of the i486 (with 8 KB internal on-chip cache) is about 3 times and 2 times higher than the 386 (with 64 KB external cache), respectively, in our laboratory benchmark tests. (See the next section for more discussion on performance). The i486 bus is also faster than the 386 bus.

(3) Consume less or equal power: The i486 consumes less power than the combination of the 386DX and the 387DX as shown in the Intel data book. The system board of the i486 also consumes less power than the 386 board as shown in the IBM presentation.

(4) Higher integration: The on-chip integration of the i486 includes a floating point unit, 8-Kbyte on-chip cache, and a memory management unit.

(5) Multiprocessor support: The i486 supports multiprocessor instructions, cache consistency protocols, second level cache, and other multiprocessor support hooks.

5.2.1.1.4 The i486 Issues And Recommendations

There are two main issues/concerns in using the i486 for the DMS at this time:

(1) The i486 has on-chip, 8K-byte internal caches for instruction and data. The cache does not have error detection and correction (EDAC) function and is not suitable for space applications. Even though the internal cache of the i486 can be disabled by software control, how much performance impact for the cache disabled on the i486 is unknown at this time.

RECOMMENDATION: It is important to measure the performance of the i486 with the on-chip cache disabled. The FII research scientists plan to implement assembly programs to disable the internal cache and find out the performance impacts in the near future. If we decide to use the i486 for the DMS, a much larger, second-level, parity-checked external cache will be necessary to improve the performance.

(2) The Work Package II is required to deliver a space-qualified processor for DMS in 1992. Intel's schedule for the i486 is: a MIL-STD-883C Class B, military-qualified i486 in the second quarter of 1992, and a Class S, space-qualified i486 in the second quarter of 1993.

RECOMMENDATION: Use the 386 to develop software because it is 100% upward binary compatible with the i486. Press Intel to have a Class S i486 available earlier than the second quarter of 1993.

5.2.1.1.5 Further Performance Enhancement

The performance of floating-point, scientific computation is important for the DMS application. The i486 has an on-chip floating-point co-processor, but the performance of the i486 can be further enhanced by adding a powerful co-processor. Weitek claims that a Weitek 4167 co-processor can have about 3 times floating-point performance improvement for the i486. An even more impressive co-processor will be the 64-bit Intel i860 co-processor that runs on a commercially available IBM Wizard Adapter.

5.2.1.1.6 The i860 Features

The Intel i860 has the following important features:

(1) A 64-bit RISC-based, high performance processor: The clock rates available for the i860 are 33 MHz and 40 MHz, and Intel claims their performance are 41 and 50 MIPS, respectively.

(2) A powerful co-processor: The i860 can be used as a stand-alone CPU, but the i860 will probably be most popular as a floating-point or graphics coprocessor for another CPU, such as the 386 or the i486.

(3) High integration: The on-chip integration of the i860 includes an integer core unit, a floating point unit, a 3-D graphics unit, 8-Kbyte data cache, 4-K byte instruction cache, and a memory management unit.

(4) MIL-STD-883C schedule: Intel plans to have the i860 meet the MIL-STD-883C Class B, military qualification in the second quarter of 1991. It takes about a year to process a space-qualified chip after its military-qualified version is available. Therefore, it can meet the 1992 delivery schedule for the Work Package II.

5.2.1.1.7 Introduction of Wizard Adapter

The Wizard Adapter card from IBM can be plugged into an IBM PS/2 Model 70 or 80 computer to improve significantly the numeric-intensive applications. The card contains an Intel i860, 2 MB of memory, and 3 IBM chips to interface with host CPU (the 386 or i486) through the IBM Micro Channel. An additional 6 MB of memory expansion is available for the Wizard card. The card can run with the IBM OS/2 operating system at this time and the IBM AIX OS later in 1990. It can be used either to run the numeric-intensive portions of 386/i486 jobs, or it can be employed to run complete jobs independently from, but under the control of, the host 386/i486. An Intel i860 Software Development Toolkit, which includes a C compiler, assembler, linker, library, etc. is available. All code required to run on the Wizard Adapter must be recompiled using the toolkit. Tasks run best on the adapter if their I/O requirements are limited. For graphics applications requiring real time updates, the Wizard Adapter allows a second dedicated display to be connected via the Micro Channel.

5.2.1.1.8 Recommendation for Wizard Adapter

The Wizard Adapter may have a great potential for the DMS. A wizard adapter, even just runs under the 386, can be estimated at about 30 MIPS, which is far superior than most of the RISC processors and the i486. Even though the Wizard Adapter will consume certain amount of power, its performance may be utilized to reduce the number of SDPs for the DMS and possibly reduce the total power consumption. Besides, multiple i860s, such as 4 processors, may be implemented on

an adapter. The host CPU (386/486) simply dispatches jobs to the shared-memory parallel processing i860's and collects the results. FII research scientists plan to measure the performance of the 386DX/486 running with a Wizard Adapter in a near future.